

PATENT APPLICATION

**CIRCUIT AND METHOD FOR PROCESSING COMMUNICATION
PACKETS AND VALID DATA BYTES**

Inventor(s): Andy P. Annadurai, a citizen of The United States, residing at
5112 Amberwood Dr.
Fremont, CA 94555

Feng Han, a citizen of The United States, residing at
1541 Poppybank Court
Pleasanton, CA 94566

Mohammed Rahman, a citizen of The United States, residing at
3689 Huff Ct.
Pleasanton, CA 94588

Chris Tsu, a citizen of The United States, residing at
13625 Surrey Lane
Saratoga, CA 95070

Assignee: SynTera Communications, Inc.
42501 Albrae Street
Suite 203
Fremont, CA 94538

Entity: Small

CIRCUIT AND METHOD FOR PROCESSING COMMUNICATION PACKETS AND VALID DATA BYTES

BACKGROUND OF THE INVENTION

5 [01] The present invention relates generally to the field of communication systems and more specifically to a method and circuit for processing data packets and valid data bytes.

10 [02] The coming of the modern information age has brought about phenomenal growth in demand for telecommunications-based products and services, driven primarily by the Internet. As the burgeoning expansion of the Internet continues along an unprecedented and unpredictable path, many new applications are foreseen and expected. These applications are placing increasing demands for ultra-high speed circuit solutions. In parallel, driven by the explosive growth in bandwidth requirements of multimedia applications, various ultra-high bit rate transmission techniques have been developed. Fiber-optic communications system speeds have increased from asynchronous-transfer-mode (ATM) rates of 155 Mb/s and synchronous optical network (SONET) rates of 622 Mb/s to the optical carrier (OC) standard of OC-192 at 10 Gb/s and OC-768 at 40 Gb/s. Over time a body of interface standards have developed to facilitate the interconnection of the various communication networks. In certain instances, however, the varying standards have presented unique challenges for system and integrated circuit designers. An example can be found in the data packet processing standards developed for SONETs.

15 [03] SONET is a transport mechanism for multiplexing data from multiple networks onto fiber optic cabling. For example, IP (Internet protocol), ATM, and T1 are among the various types of networks that are interconnected via SONET rings. Since SONET interconnects different network types, it employs routers for converting protocols from one network type to another. In the case of IP protocols, for example, routers are used for converting IP packets for transmission on the SONET ring.

20 [04] A typical router contains a line card for receiving data packets on one end, performing necessary conversions and sending out the packets at the other end. Among other components, line cards include a framer for framing/de-framing data packets, and a processor for performing protocol conversion and for controlling packet traffic. The framer communicates with the processor through an interface known as the SPI-4 (system packet interface). The SPI-4 interface is a standard defined by a consortium of communication

companies. This standard defines packet and cell transfer standards between a physical layer device (i.e., the framer) and a link layer device (i.e., the processor). The SPI-4 interface requires 16-bit data processing. That is, data is processed in 16 bit words, containing two bytes of data. Each packet contains an even number of bytes, each of which must be valid.

5 However, many functional blocks employed in telecommunication systems including, for example, most network processors process data on byte (8 bit) boundaries as opposed to word (16 bit) boundaries. In such systems, a packet need only contain byte multiples for proper processing. For instance, it may contain odd multiples such as three bytes, seven bytes etc. Or, it may contain even multiples such as two bytes, eight bytes, etc. So long as the packet
10 contains multiple bytes, it is properly processed. Upon arriving at the SPI-4 interface, however, data bytes contained in a packet must be processed as 16 bit words. The data bytes are first grouped into 16 bit words (two bytes) for processing. This presents no problems so long the packet contains an even number of bytes. If, however, the number of bytes is odd, a single byte will remain after grouping. For example, for a three byte packet, the first two
15 bytes are grouped into a 16 bit word. The last byte can only be grouped with an empty byte. This empty byte is of course invalid and should be processed as such.

[05] Identifying such invalid bytes and processing them in a manner that does not compromise the integrity of the data and speed of transmission is a challenging task. There are other similar incompatibilities between the varying interface standards and
20 established infrastructure that give rise to such challenges.

[06] There is, therefore, a need for methods and systems that can resolve such incompatibilities and the problems that may arise from erroneous processing of data packets in communication networks such as SONETs.

25 BRIEF SUMMARY OF THE INVENTION

[07] The present invention provides methods and systems for improved processing of data packets in communication systems such as a synchronous optical network (SONET). Data packets received by communication systems often contain invalid data bytes. The present invention provides, in one embodiment, a method for processing data that drops
30 these invalid bytes, and concatenates valid bytes for further processing. Concatenation refers to linking of data bytes across packet boundaries for further packaging into high speed packets.

[08] According to a first aspect of the present invention, a method for processing data in a communication system, wherein the method includes the step of

receiving a first data packet that comprises a plurality of pairs of bytes of data. Thereafter, the first data packet is examined to determine if the final pair of bytes comprises a valid byte and an invalid byte. If an invalid byte is present, the valid byte is concatenated with a byte from a subsequent packet. Concatenation is initiated by dropping the invalid byte. This leaves the valid byte in the first data packet. Thereafter, the method includes the steps of receiving a second data packet, and determining whether the second packet begins with a valid byte. If so, this valid byte is concatenated with the valid byte of the first data packet byte. In this manner, the first packet is completed with valid bytes, and may be sent out or packaged into a super packet. In one embodiment, data bytes within the second packet are shifted to occupy empty byte locations created by the preceding bytes.

[09] According to another aspect of the present invention, a concatenation circuitry for processing data packets is disclosed. The concatenation circuitry features logic circuits for determining that a first packet ends with a valid byte and an invalid byte of data, and logic circuitry for dropping the invalid byte of data. Further, the concatenation circuitry includes logic circuitry for determining that a second data packet begins with a valid byte of data, and logic circuitry for concatenating the valid byte of the first data packet with the valid byte of the second data packet.

[10] A further understanding of the nature and advantages of the present invention may be realized by reference to the remaining portions of the specification and the attached drawings. References to "steps" of the present invention should not be construed as limited to "step plus function" means, and is not intended to refer to a specific order for implementing the invention. In the drawings, the same reference numbers indicate identical or functionally similar elements.

BRIEF DESCRIPTION OF THE DRAWINGS

[11] Figure 1A is an exemplary communication system in which the present invention is implemented.

[12] Figure 1B illustrates concatenating data bytes from two packets both packets ending with invalid bytes, in accordance with an exemplary embodiment of the present invention.

[13] Figure 2 illustrates concatenating data bytes from three or more packets in accordance with an exemplary embodiment of the present invention.

[14] Figure 3 illustrates an exemplary circuitry for outputting 16 bit words when no byte shifting is required in accordance with an embodiment of the present invention.

[15] Figure 4 illustrates a circuitry for shifting data bytes to produce exemplary 16 bit data in accordance with an embodiment of the present invention.

[16] Figure 5 illustrates exemplary circuitry for outputting 16 bit data when prior packets end with even valid bytes and the current packets ends with a single valid byte.

[17] Figure 6 illustrates exemplary circuitry for dropping invalid bytes, concatenating and shifting data bytes in accordance with an exemplary embodiment of the present invention.

[18] Figure 7 is an exemplary state diagram for implementing control logic of the byte multiplexing of Figure 6.

DETAILED DESCRIPTION OF THE INVENTION

[19] Figure 1A is an exemplary communication system 100 in which the present invention is implemented.

[20] Communication system 100 may be a transport mechanism for delivering high-speed data to a destination 132 via fiber optic cabling 130. An example of such a network is SONET. Various specific embodiments of the present invention are described herein in the context of a SONET implemented using 128 bit wide bus. It is to be understood, however, that the invention is applicable to other types of communication systems and networks, that specific bus width or other implementation-specific values and numbers provided herein are for illustrative purposes only, and that the invention applies to telecommunication systems with other implementations.

[21] Among other components, communication system 100 includes fiber optic cabling 130 for receiving (and transmitting) high speed data such as IP packets. System 100 further includes a serializer-deserializer (or SERDES) 134 for serializing and de-serializing data packets, and a framer 136 for decoding data packets, and for forwarding the payload to destination 132. In addition, a network processor 138 receives the deserialized data and outputs the payload along a byte boundary (8 bit) to destination 132.

[22] For the receive direction, at interface 140 between SERDES 134 and framer 136, data bytes are processed along 16 bit boundaries at interface 140, and thereafter, grouped into 128 bit packets in framer 136. Following this, at interface 142 between framer 136 and processor 138, which is the SPI-4 interface, data bytes are processed again along 16 bit boundaries. These bytes are then output along eight byte boundaries by processor 138. Similarly, for the transmit direction, a reverse process may occur from destination 132 to network side 144. Thus, at the SPI-4 interface, packets containing eight bit bytes (or

10687228.022702
multiples thereof) are received from processor 138. These are processed and supplied to framer 136 as 16 bit words. At framer 136, eight groups of 16 bit words are grouped as 128 bit packets, etc. The words are packed into a 128 bit wide data bus and sent to a transmit first-in first-out (FIFO) structure (not shown) to be processed further by framer 136.

5 [23] As previously noted, there is a chance that the last 16 bit word of packets received by framer 136 may contain invalid data bytes. This may happen when such packets contain odd number of bytes. In accordance with the present invention, exemplary 16 bit words received by framer 136 are processed to identify invalid bytes. Once identified, the invalid bytes are dropped, so that valid bytes are grouped into 128 bit packets.

10 [24] Although the process occurs at the SPI-4 interface, one of ordinary skill in the art will realize that this is for illustrative purposes only, and that the process according to the invention may be applied to other instances where an invalid byte has been detected. It is advantageous to drop the invalid byte at the SPI-4 interface, before such a byte becomes further processed along the data path. After the invalid byte is dropped, a valid
15 byte from one data packet is concatenated with a valid byte from the subsequent data packet. As used herein "concatenation" refers to the linking of data bytes from different packets into a high rate envelope.

20 [25] Figure 1B illustrates concatenating data bytes from two exemplary packets 101 and 102 both packets ending with invalid bytes 103 and 104, in accordance with an exemplary embodiment of the present invention.

[26] More specifically, according to this exemplary embodiment of the invention, invalid byte 103 is dropped in order to concatenate valid data byte 105 from packet 101 with valid data byte 106 from packet 102. As shown, all bytes of packet 101 are valid except for invalid byte 103.

25 [27] The concatenation process is initiated by first dropping invalid byte 103. In this fashion, an empty byte is created for valid byte 106 of packet 102. Next, valid byte 106 is concatenated with valid byte 105 to provide the last word of packet 101. This concatenation completes packet 101. At this point, packet 101 may be transmitted or packed into a high speed super packet. In one embodiment, this super packet contains five 16 bit
30 packets for a total of 128 bits.

[28] While the concatenation process as described thus far completes packet 101, it includes a byte shifting process to complete succeeding packet 102. This is because valid byte 106 is now empty, its content having been transferred to packet 101. To fill byte 106, byte 107 is shifted to byte 106. Similarly, byte 108 is shifted to byte 107 This shifting

continues until byte 109 is shifted to byte 110. At this point, no further shifting occurs since the last remaining word contains only one byte 104 that is invalid. The byte shifting process is thus terminated. In this manner, only valid bytes are contained within packet 102 in accordance with this embodiment of the present invention. It should be noted that after byte 109 is shifted, the entire last word of packet 102 becomes empty. For this reason, when packet 102 is transmitted, the receiving system is informed that the last word is empty and can be ignored.

[29] Figure 2 illustrates concatenating data bytes from three or more packets 210, 212, N, in accordance with another exemplary embodiment of the present invention.

[30] In this specific embodiment, concatenation is illustrated where second packet 212 ends with a valid data byte 215 so that concatenation continues until a packet N ending with an invalid byte 213 is received. In Figure 2, the concatenation process first drops invalid byte 211 of packet 210. Next, valid byte 219 is concatenated with valid byte 205 to complete the last word of packet 210. Next, byte shifting starts for the second packet 212. Byte shifting occurs in a similar manner as described with reference to Figure 1B, except that because packet 212 ends with valid data 215, byte shifting continues after packet 212. That is, after byte 214 is shifted to byte 216, byte 215 is shifted to byte 214. Thus byte 214 has valid data while byte 215 becomes empty. Packet 212 awaits the arrival of a subsequent packet having a valid byte. The first valid byte of this subsequent packet is shifted to byte 215 to complete the last word of packet 212. Thereafter, packet 212 is sent out.

[31] Byte shifting continues for subsequent (N-1) packets all of which end with two valid bytes. Byte shifting ends when packet N is reached which ends with one valid byte 217 and one invalid byte 213. At this point, valid byte 217 is shifted to byte 218, and invalid byte 213 is dropped. Byte shifting then comes to an end.

[32] Figures 1B and 2 show two different conditions for byte treatment according to the present invention.

[33] There are, however, a number of other conditions that are also addressed by the method of the present invention. Every possible condition for different byte treatment inside packets and at packet boundaries can be defined by any combination of the following:

[34] There are two possible conditions (1) and (2), below, for byte treatment inside a packet. Note that these two conditions do not address packet boundary

conditions (i.e., end or start of a packet). There are four possible packet boundary conditions, (3) through (6), below:

[35] (1) all prior packets carry even number of valid bytes of data, and therefore no byte shifting occurs inside a packet;

[36] (2) one prior packet ends with a single valid byte, and therefore byte shifting occurs inside the subsequent packets until a packet ending with a single valid byte arrives;

[37] (3) all previous packets end with an even number of valid bytes, and the current packet ends with two valid bytes;

[38] (4) all previous packets end with an even number of valid bytes, and the current packet ends with only one valid byte;

[39] (5) a previous packet ends with a single valid byte, and byte shifting is occurring in the current packet, and the current packet ends with two valid bytes; and

[40] (6) a previous packet ended with a single valid byte, byte shifting is occurring in the current packet, and the current packet ends with one valid byte.

[41] Various exemplary circuits for implementing the functionality required by the above six different conditions for byte treatment will be presented hereinafter.

Referring to Figure 3, there is shown an exemplary circuitry 300 for generating 16-bit words when no byte shifting is required in accordance with an embodiment of the present invention. Specifically, circuitry 300 implements the functionality associated with conditions (1) and (3) above, where neither byte dropping nor byte shifting is required. For these situations, no packet ends with a valid byte.

[42] In this exemplary implementation, among other components, circuitry 300 includes two levels of registers. A first level includes register 306 and register 308, both of which are 8-bit registers. Register 306 has an input 300 for receiving bits 8 through 15 of an incoming data word. This is represented as TDAT[15:8]. Register 308 has an input 303 for receiving bits 0 through 7 of the incoming data word. This input is represented as TDAT[7:0]. The second level of registers also includes two 8-bit registers 307 and 309. The input of register 307 is coupled to output 301 of register 306. Output 302 of register 307 provides data output for the 16-bit word. This is shown as DataOut[15:8]. The input of register 309 is connected to output 304 of register 308. This is represented as DataReg[7:0]. The output 305 of data register 309 is represented as DataOut[7:0].

[43] In operation, output 302 is a copy of input 300 that is two clock cycles delayed. Similarly, output 305 is a copy of input 303 but two clock cycles delayed. When

received, the input data is sent through circuitry 300 with no shifts. Specifically, bit locations for input data are maintained from input to output. For example, referring to Fig. 1B, byte 122 with bits 15:8 is passed through registers 306 and 307 without changing bit locations. On the same cycle, byte 124 is passed through registers 308 and 309 without bit location changes. On the next clock cycle, the next word is passed through without changes, and so forth. In this manner, an exemplary circuit is provided for passing received data bytes without changes, when such bytes are valid.

[44] Figure 4 illustrates a circuitry 400 for shifting data bytes to generate exemplary 16-bit data words in accordance with another embodiment of the present invention. Specifically, circuitry 400 provides data paths for conditions (2) and (5) above, where a prior packet has ended with a single valid byte, and subsequent packets, including the current packet, end with two valid bytes. Thus, this condition performs essentially a continuous byte shifting operation. Among other components, circuitry 400 includes two levels of registers as in Figure 3. These registers are registers 406, 408, 407 and 409. As shown, register 406 comprises an input 400 for TDAT[15:8], and an output 401 for DataReg[15:8]. Output 401 is coupled to the input of register 407. Further, output 401 is coupled to output 405, DataOut[7:0]. Register 408 has an input 403 for TDAT[7:0] and an output 404 for DataReg[7:0]. Output 404 is coupled to the input of register 409 having an output 402 for DataOut[15:8].

[45] The operation of circuitry 400 will be described with reference to Fig. 2. Packet 210 can be considered as the prior packet ending with one valid byte (205), while packet 212 is the first one of several succeeding packets each ending with two valid bytes (214 and 215 in the case of packet 212). Accordingly, circuit 400 operates to continuously shift data, one byte at a time, within a word, e.g., 215 to 214, between words in the same packet, e.g., 214 to 216, and between words falling on the boundary between two packets ending with valid bytes, e.g., 224 to 215.

[46] Referring to Fig. 4 and considering, for example, the boundary condition between packets 212 and 222, during a first clock cycle, input 400 TDAT[15:8] carries values for byte 214 (Fig. 2), and input 403 TDAT[7:0] carries invalid byte 215. These 16 bits arrive and are stored in registers 406 and 408, in one clock cycle. On the next clock cycle, bytes 214 and 215 are clocked into the next level of registers 407 and 409, respectively, and bytes 224 [15:8] and 226 [7:0] arrive at input terminals 400 and 403, respectively, and are clocked into registers 406 and 408, respectively.

[47] Note that the data in registers 407 and 409 are one clock period delayed with respect to data in registers 406 and 408. As shown in Fig. 4, the 16-bit output DataOut[15:0] is formed by using byte 215 at the output of register 409 for bits DataOut[15:8], and using byte 224 at the output of register 406 for bits DataOut[7:0]. This results in shifting byte 215 as shown by arrow 228 in Fig. 2, and shifting byte 224 as shown by arrow 230 in Fig. 2. Circuit 400 thus implements the continuous byte shifting in the same manner as shown by the arrows in Fig. 2

[48] Figure 5 illustrates exemplary circuitry 500 for generating 16-bit data for the condition where prior packets end with even valid bytes and the current packet ends with a single valid bytes. Specifically, circuitry 500 is designed to address boundary condition (4) listed above. Among other components, circuitry 500 includes two levels of registers 506 and 508 that process byte [15:8] of each word, and registers 507 and 509 that process byte [7:0] of each word. As shown, register 506 comprises an input 500 for TDAT[15:8], and an output 501 for DataReg[15:8]. Output 501 is coupled to the input of register 507, and is coupled to output 505, DataOut[7:0]. Register 507 has an output 502 for DataReg[15:8]. Register 508 has an input 503 for TDAT[7:0] and an output 504 for DataReg[7:0]. Output 504 is coupled to the input of register 509.

[49] The operation of circuit 500 can be described with reference to the exemplary packets shown in Figure 1B, where packet 101 is considered the current packet ending with one invalid byte. In a similar manner as described in connection with the operation of circuit 400 above, at the packet boundary, two consecutive clock cycles cause registers 507 and 509 to store bytes 105 and 103, respectively, and registers 506 and 508 to store bytes 106 and 107, respectively. To form the output data DataOut[15:0], circuit 500 taps output 502 of register 507, or byte 105, for DataOut[15:8], and taps output 501 of register 506, or byte 106, for DataOut[7:0], as shown. As shown in Fig. 5, the output of register 509, which in this condition stores invalid byte 103 is not used. This results in invalid byte 103 being dropped, and bits [15:8] of byte 105 begin concatenated with bits [15:8] of byte 106, although they are one clock cycle apart.

[50] The final condition described in (6) above addresses the situation where a prior packet ends with a single valid byte, and byte shifting is occurring in a current packet that also ends with a single valid byte. An example of this condition is illustrated in Fig. 1B when considering packet 102 as the current packet. Due to the byte shifting that is initiated when a first packet is encountered that ends with one valid byte, the result of a second packet ending with one valid byte is that, after byte shifting in the second packet, it

ends with a single invalid byte in its last word. This last invalid word can thus be dropped by skipping one clock cycle.

Figure 6 illustrates an exemplary circuitry 600 for implementing the packet treatment as described above under condition (6), according to one embodiment of the invention. In this specific example, since packet condition (6) is essentially a combination of the other conditions, circuitry 600 is also a combination of circuitry 400 of Fig. 4 and circuitry 500 of Fig. 5. Among other components, circuitry 600 includes a register 607 having an input port 601 for receiving input data TDAT[15:8] and having an output port 603 coupled to the input port of an 8-bit 2:1 multiplexer (MUX) 609. A register 608 has an input port 602 for receiving input data TDAT[7:0], and an output port 604 that couples to another input of MUX 609 as well as an input of another 8-bit 2:1 MUX 610. The other input of MUX 610 receives input data TDAT[15:8].

MUX 609 receives a control signal on line 642 and MUX 610 receives its control signal on line 656. Further, circuitry 600 includes a register 611 having an input port coupled to output port 640 of MUX 609, and a register 612 having an input port coupled to output port 650 of MUX 610. Register 611 generates output data dout[15:8], and register 612 generates output data dout[7:0]. Further, TDAT[15:8] is received from an output of register 656 while TDAT[0:7] is received from an output of register 658.

[51] In operation, for packet conditions (1) and (3) outlined above, the control signal on line 642 for MUX 609 selects the output of register 607, dataReg[15:0] on line 603, to pass to register 611, and the control signal on line 656 for MUX 610 selects the output of register 608, dataReg[7:0] on line 604, to pass to register 612. This selection results in a circuit equivalent to that of Figure 3, which handles the condition where there is no need for dropping or shifting any bytes.

[52] For packet conditions (2) and (5), the control signal on line 642 for MUX 640 selects the output of register 608, dataReg[7:0] on line 604, to pass to register 611, and the control signal on line 656 for MUX 610 selects input data TDAT[15:8] to pass to register 612. This selection results in a circuit equivalent to that of Figure 4, which implements the required continuous byte shifting.

[53] For packet condition (4), the MUX control signals cause MUX 609 to select dataReg[15:8] on line 603 to pass to register 611, and cause MUX 656 to select TDAT[15:8] to pass to register 612. This selection results in a circuit equivalent to that of Figure 5 implementing the required functionality for packet condition (4). In this manner,

MUXes 609 and 610 allow circuit 600 to be reconfigured the data path to address all of the different combinations of byte treatment.

[54] Figure 7 is an exemplary embodiment of a state machine 700 for providing control signals to selection line 642 of MUX 609 and selection line 656 of MUX 610. It is to be understood that various logic circuit implementations for the state diagram of Fig. 7 are possible based on known logic design techniques. As shown, four states represent combinations of the six conditions discussed with reference to Figs. 1B and 2. For example, EVEN state 710 represents a combination of conditions (1) and (3). In EVEN state 710, no odd numbered packet is received. Thus, there is no need to do any byte shifting. The transition from EVEN state 710 to EVEN2ODD state 712 takes place when the control word shows that the current packet ends with a single valid byte (TCTL & TDAT[14] & TDAT[13] 718). EVEN2ODD state 712 lasts for only one clock cycle, during which time, the last valid byte of the current packet is concatenated with the first byte of the following packet.

[55] Thereafter, state machine 700 enters the ODD state 714, in which all bytes are left shifted by 1 byte position. State machine 700 comes out of ODD state 714 when the control word shows that a second packet ending with a single valid byte has arrived (TCTL & TDAT[14] & TDAT[13], 722). This is when ODD2EVEN state 716 is entered. ODD2EVEN state 716 lasts for only 1 clock cycle. In this state, a control signal is sent out to the transmit FIFO, informing the FIFO to drop one whole word from the incoming transmit data. After ODD2EVEN state 716, state machine 700 returns to EVEN state 710.

[56] In this fashion, the present invention provides a circuit and method for processing communication packets and valid data bytes in a communication system. While the above is a complete description of exemplary specific embodiments of the invention, additional embodiments are also possible.

[57] Therefore, the above description should not be taken as limiting the scope of the invention, which is defined by the appended claims along with their full scope of equivalents.